Load Selection and Home-Return Policies in Historical Truckload Replay

A comparative development study of five terminals and two driver schedules

Abstract

Truckload dispatchers must select loads while balancing near-term contribution, time, equipment and hours-of-service feasibility, future positioning, and an eventual return home. We compare 12 implemented selection policies using historical load-board replay for five terminals, Monday first pickups, and two declared duty cycles. The development grid contains 1,140 completed policy–terminal–start–schedule replays: 600 for a six-day-out, one-day-home schedule (6/1) and 540 for a twelve-day-out, two-day-home schedule (12/2). A common carrier cost profile is used and cycle contribution is projected to a year using fixed seven- or fourteen-day denominators. Across 50 terminal-start observations per policy, pair-preferred selection with always-on home valuation had the highest 6/1 mean annualized contribution, $236; across 45 observations per policy, the four-day home-valuation policy had the highest 12/2 mean, $13,630. Relative to pair-preferred open selection on matched terminal-start cases, the corresponding mean changes were $403.76 per 6/1 cycle and $237.39 per 12/2 cycle; their medians were $396.01 and $7.68. No selector had a positive weekly mean or median in every tested week. Missing dispatch-day captures were handled using earlier boards in 490 of 1,140 cases, an assumption that cannot establish subsequent offer availability. These results describe an in-sample simulator and are not estimates of realized earnings or proof of policy superiority in operation.

Keywords: freight transportation; spot market; truckload dispatch; load selection; home time; historical replay; contribution margin

1. Introduction

A dispatcher comparing two loads faces more than two prices. A higher-paying load can require more unpaid travel, occupy the driver longer, deliver into a weaker market, or make the required trip home expensive. A shorter or less lucrative initial load may allow the truck to accept another

booking sooner. Conversely, a high-value load that requires waiting can still be the better choice when its additional contribution compensates for the elapsed time. Real dispatch decisions also depend on broker reliability, appointment confirmation, customer requirements, driver preferences, negotiation, and whether an advertised load remains available. A numerical selector can support these decisions only to the extent that those features are represented in its inputs and model.

We ask three questions. First, how do implemented immediate-load and limited-lookahead rules compare when assessed on a common simulated cycle contribution? Second, how do policies that include an estimated trip home in load selection behave across different home-time schedules? Third, how stable are outcomes across terminal locations and tested start weeks? The unit of decision is a single truck, not a fleet assignment. The experiment is a development study: policies were explored and tuned using the same historical archive subsequently reported here.

Our main contribution is a transparent comparison of complete simulated cycle outcomes, including unfavorable and zero-load outcomes, against a common pair-preferred open-route policy. We also distinguish the local score used to select a load from the fixed-cycle annualization used for evaluation. We do not claim the tested settings are optimal, that offers were bookable, or that 1,140 replays constitute 1,140 independent market observations.

2. Related research

Online freight routing has been studied as a sequence of choices under uncertain future opportunities. Miller, Nie, and Liu [1] study hyperpath routing in an online freight exchange. The present work instead evaluates implemented decision rules on captured load-board observations and does not reproduce their model or estimate its performance. Dong et al. [2] formulate the selection of pickup-and-delivery requests during a return to a depot under capacity and required return-time constraints. Their backhaul problem is related to our homeward decision, but our policies repeatedly decide in a rolling replay rather than solve their formulation.

Dumouchelle, Frejinger, and Lodi [3] study sequential booking control with capacity reservation and operational fulfillment costs; their applications are distribution and air cargo, not this truckload setting. Chandrasekaran [4] presents a public-calibrated simulated truckload benchmark with explicit operational feasibility, economic assumptions, and decision-latency considerations. Ojha et al. [5] examine load planning in parcel service networks and emphasize

consistency and computational support for human planners. These studies motivate explicit state, objective, feasibility, and reproducibility descriptions; their reported outcomes do not validate the policies evaluated here.

2.1. Freight market structure and economics

The trucking market includes self-driving owner-operators, carriers employing multiple drivers, large asset-based fleets, and intermediaries that arrange transportation using other carriers. These categories overlap: an owner-operator may run under another carrier's authority or hold authority independently; a fleet-size statistic does not identify who drives or owns each truck. Nor is there one universal boundary between a small and a large carrier. The fleet-size categories below use tractor counts solely to make the available registration data interpretable.

According to American Trucking Associations (ATA), citing U.S. Department of Transportation data, almost 580,000 active U.S. motor carriers owned or leased at least one tractor as of June 2025. ATA reports that 91.5% operated ten or fewer trucks and 99.3% operated 100 or fewer [6]. Taking differences between the rounded percentages gives approximately 7.8% with 11–100 trucks and 0.7% with more than 100. Applying those proportions to the approximate total yields roughly 531,000, 45,000, and 4,000 carriers, respectively; these are arithmetic illustrations from rounded figures, not precise counts. They are not counts of owner-operators, long-haul spot carriers, or all trucking firms, and they do not establish the share of trucks, loads, or revenue handled by each category. ATA estimates $906 billion in gross trucking freight revenue for 2024 primary shipments [6]; that broad trucking total is not an estimate of the spot truckload segment used in this replay.

Table A. Illustrative fleet-size groupings calculated from ATA's June 2025 registered-carrier summary [6]. The three categories are mutually exclusive, but the owner-operator business model can appear within more than one category.

Tractors per registered carrier	Share of carriers	Approximate carriers	What the count does not establish
1–10	91.5%	531,000	Number of owner-operators, their load counts, or revenue
11–100	7.8%	45,000	A representative small-carrier margin or revenue per truck

Tractors per registered carrier	Share of carriers	Approximate carriers	What the count does not establish
More than 100	0.7%	4,000	The group’s share of loads or the entire freight bill

The following 2025 company segment figures illustrate scale and reported operating economics. Revenue is segment revenue, not the total dollar value of shippers’ goods. Operating margin here is segment operating income divided by segment revenue; it is not net income margin. The last column is our calculation of segment operating income divided by average tractors, not a company-reported contribution per truck. Tractors include company and independent-contractor equipment where the filings specify. Segment mixes and accounting boundaries differ.

Table B. Selected reported 2025 carrier segment results and calculated income per average tractor [7–9]. Rounded values are not matched to the replay’s cost boundary.

Organization and segment	Segment revenue	Average tractors	Operating income	Operating margin	Income per tractor-year
Werner, Truckload Transportation Services	$2.052 billion	7,437	$16.426 million	0.80%	$2,209
Knight-Swift, truckload	$4.865 billion	21,428	$147.232 million	3.03%	$6,871
Marten, truckload	$421.729 million	1,668	$0.825 million	0.20%	$495

These are single-year segment results, not typical margins for all large fleets. Werner’s Truckload Transportation Services segment includes dedicated and one-way services and some non-trucking revenue; the 2025 result was affected by unusual items. Knight-Swift’s truckload segment includes several service types, and Marten’s truckload operations include company and contractor tractors [7–9]. Dividing aggregate income by average tractors is not an individual driver’s pay and may also reflect segment revenue and expenses that do not map one-for-one to an individual truck.

For owner-operators, ATBS reports 2025 average net income of $71,808 for its client population and approximately 95,000 annual miles for the average independent contractor in its discussion [10]. This is an owner-operator sample, not a representative census; its net income includes a return to the owner’s own driving labor, and the precise expense and tax boundaries require reconciliation. It should not be labeled a carrier operating margin, a driver wage, or a direct

comparison with Table B. A 2025 per-owner load count, gross revenue, and profit margin were not established from the accessible source, and a reliable comparable load-count and revenue distribution for 1–10- or 11–100-truck carrier businesses was not established for this paper.

Landstar offers another relevant, but differently defined, view of owner-supplied capacity. In fiscal 2025, its business capacity owner (BCO) independent contractors hauled 798,050 loads that generated $1.804 billion of Landstar-reported revenue, or $2,260 per load; Landstar states that fuel surcharges billed to customers on BCO-generated revenue are excluded from that revenue. At fiscal year-end, BCOs provided 8,514 trucks, approximately 96% of which were provided by BCOs with five or fewer trucks in Landstar's network [11]. The truck count is a year-end snapshot, not an annual average; dividing annual loads or revenue by it would not establish annual productivity per truck. Landstar revenue on BCO loads is not the BCO's take-home revenue or earnings.

Landstar also reports non-GAAP variable contribution, defined as company revenue less variable costs of revenue, and the corresponding variable contribution margin. For fiscal 2025, these were $668.020 million and 14.1% for the company overall [11]. This includes business beyond the BCO network; the company-wide amount must not be allocated to BCO trucks as though it were a BCO-only numerator. Landstar describes BCO-related contribution margin as generally a fixed percentage because of its contractor and agent contracts, but that company margin is not the contractor's revenue after that contractor's fuel, maintenance, equipment and other costs. We did not identify a directly reported BCO-specific contribution dollar amount per truck in the reviewed filing. Landstar's contribution measure and this paper's modeled carrier contribution therefore have different cost and organizational boundaries.

Load count measures completed or arranged loads only when the source says so; it is not the number of board postings. As a disclosed example, Marten reported 95,951 loads and $149.826 million in 2025 brokerage revenue, equivalent to approximately $1,562 of brokerage revenue per reported load [9]. That is a brokerage-segment ratio, not the price received by the hauling carrier, and it must not be presented as the load count or average rate for Marten's asset-based truckload tractors. Its truckload segment separately reported 153.699 million total miles and a 520-mile average trip, but those figures alone do not establish the number of truckload shipments because total miles and average trip length do not necessarily have matching definitions.

The volume arithmetic is straightforward but should not be mistaken for empirical averages. For example, a hypothetical one-truck operation completing 120 loads at $1,500 of gross revenue each would generate $180,000 in load revenue; ten identical trucks would generate 1,200 loads and $1.8 million. Actual loads per truck, revenue per load, fixed costs, driver compensation, and

utilization vary. No profit margin follows from these assumptions without an explicit cost model. Similarly, an operating ratio of 97% means operating expenses are 97% of revenue on that measure, leaving a 3% operating margin; it does not imply a 3% net margin after interest and taxes. The narrow margins in the selected 2025 filings make cost-boundary alignment important when interpreting the much less complete simulated contribution in this paper.

3. Decision setting and evaluation measures

At a simulated decision time the truck has a location, clock, accumulated duty state, and a required home deadline. The replay reveals a selected board snapshot, filters candidate loads through the implemented load and timing rules, and evaluates feasible transitions. A load is represented by a posted price, pickup and delivery locations and windows, and loaded mileage. Deadhead to pickup and return-home movement are modeled separately. The simulator applies its hours-of-service and home-return checks; these are projections based on available inputs, not proof of legal compliance or of actual appointment performance.

Load contribution is modeled as posted revenue less incremental modeled cost. The carrier profile assumes approximately $1.897 per total mile, including fuel, maintenance, tires, tolls, employee driver pay, and benefits, with zero commission. These are fixed model inputs, not observed costs for the sampled postings. The per-cycle ledger includes modeled loaded, pickup-deadhead, and homeward costs. Other overhead and business-specific expenses are not established as fully covered, so contribution is not claimed to be operating income or accounting net profit. Posted prices are treated as model revenues without evidence of negotiation or acceptance.

Annualized contribution A is defined as $A = C \times 365 / D$, where C is simulated contribution for the entire declared cycle and D is seven calendar days for 6/1 or fourteen for 12/2. The denominator does not shorten when a driver returns home early; a zero-load, on-time home outcome contributes zero. A is a repeated-cycle projection, not annual realized profit. During load selection, the one- and two-load annualized-contribution policies instead evaluate a candidate bundle relative to elapsed committed time; their local ranking score is not the same quantity as A. Waiting affects that elapsed-time score, but the selected load can still justify waiting when its contribution is sufficiently high.

The all-mile rule directly maximizes posted price divided by pickup deadhead plus loaded miles. Under the constant cost per mile and zero-commission assumptions, this has the same ranking as load contribution per those miles. We label it “contribution per all mile (gross-RPM

implementation)” for interpretation; the implementation itself scores gross RPM. The denominator does not include later repositioning, the final trip home, or the time consumed by the load. With load-specific cost rates the two rankings need not agree.

4. Data, policies, and experimental design

The historical archive covers a limited period in early 2025. The five starting terminals are Atlanta, Chicago, Columbus, Nashville, and Raleigh. First pickup is required on the assigned Monday where a load is obtained. Ten Mondays from February 3 through April 7 are eligible for 6/1; nine from February 3 through March 31 are eligible for 12/2. The April 7 12/2 deadline and both April 14 deadlines fall outside the observed archive horizon. The comparison therefore has 50 terminal-week scenarios for each 6/1 policy and 45 for each 12/2 policy. No-load outcomes are retained; they do not have an actual pickup date. The analysis does not establish national representativeness, seasonal variation, owner-operator economics, or outcomes for untested terminal locations.

The common study eligibility screen requires positive posted price and at least 250 reported loaded miles, an over-the-road proxy that excludes shorter opportunities. A mileage-consistency screen also rejects a load when the estimated direct distance exceeds 50 miles and reported loaded mileage is less than 90% of that estimate. Equipment and implemented timing and feasibility checks further restrict candidates. These screens shape the choice set, and the mileage proxy is not a verified trip distance. At a booking decision, the two-load candidate paths are assembled from the single revealed board rather than a future board; the replay assumes the selected offers can be booked together. It does not establish that a dispatcher could secure both offers in practice.

The 12 policy configurations are summarized in Table 1. The pair-preferred rule uses a feasible two-load bundle when one is available and otherwise falls back to one load. Flexible selection ranks eligible one- and two-load candidates together. For local bundle ranking, elapsed time starts at projected driver availability and ends at the last paid delivery; home-valued scoring includes the estimated home trip and its time. The implementation applies a one-hour minimum to elapsed duration for open-bundle annualization. The open objective evaluates the candidate bundle without adding the terminal home leg to the ranking score, although a feasible return is required. The always-on home-valued variant includes an estimated immediate home return in

the candidate objective. The four-day variant changes to home-valued ranking when projected availability is within four elapsed days of the deadline; the initial booking retains the open objective. These are alternative policies, not clean measurements of an isolated return leg.

Table 1. Evaluated decision rules. Every policy is evaluated on the same cycle accounting basis, although the selection algorithms and their candidate construction are not identical. In these names, ?two loads? means a feasible pair is preferred; the rule falls back to one load when a pair is unavailable. ?Per time? refers to modeled contribution divided by committed elapsed time, not contribution per mile.

Policy	Primary selection principle
Contribution per time: prefer two loads	Prefer a feasible two-load bundle; rank using bundle contribution over committed time without a home-cost adjustment.
Contribution per time: one load	Rank a single-load candidate using contribution over committed time.
Contribution per time: best one or two loads	Compare eligible one- and two-load candidates using contribution over committed time.
Two loads: include return-home cost and time every time	Pair-preferred selection with estimated immediate-home value in the objective.
Two loads: include return-home cost and time in final 4 days	Switch from open to home-valued ranking near the deadline, subject to the initial-booking exception.
Highest load revenue	Favor higher posted revenue.
Highest revenue per loaded mile	Favor higher revenue per loaded mile.
Shortest empty drive to pickup	Favor lower pickup deadhead.
Highest contribution on one load	Favor modeled incremental contribution.
Highest contribution per total load mile	Favor gross revenue per pickup-deadhead plus loaded mile; equivalent ranking to modeled contribution per those miles under this profile.
Meet dispatch thresholds, then highest contribution	Apply stated thresholds, then favor incremental contribution.
Contribution plus average destination value	Add an average-based destination prior to immediate contribution.

The selector names above correspond respectively to the implementation labels pair-preferred open, one-load open, flexible open, pair-preferred always-on return, pair-preferred four-day return, gross revenue, loaded revenue per mile, minimum deadhead, immediate contribution, all-mile revenue per mile, threshold dispatcher, and legacy backhaul aware.

The threshold rule requires at least $1.50 revenue per loaded mile and nonnegative modeled incremental contribution, at most 150 miles of pickup deadhead, and at most 36 hours of pickup wait. The legacy backhaul rule is an average-based benchmark, not the separately considered top-5% destination-opportunity rule. Its ranking has not been fully independently reconstructed, and its results require particular caution. The excluded contribution-per-day and top-5% destination selectors are not part of this completed matrix. The baseline rules check simulated feasibility, including the ability to return home; a common feasibility screen does not imply identical candidate sets or equal search effort across algorithms.

When the expected dispatch-day capture is absent, the enabled missing-dispatch setting uses an earlier available board. The fallback uses retrospective knowledge of capture coverage, and a prior listing is not evidence that the same offer could actually have been booked later. Across the selected attempts, 559 of 4,733 recorded optimization requests used fallback, in 490 of 1,140 cycles; no no-board events were recorded. Initial booking used the earlier Thursday board when the natural Friday capture was missing. The relevant missing target business dates were January 31, February 27, March 4, March 7, and March 20. Coverage and booking assumptions must therefore be considered when comparing policies, especially when a selector appears to benefit from a specific week.

5. Results

All 1,140 selected cases completed with verified recorded identities and on-time modeled home arrival. The 1,110 cases with home arrival after load activity and 30 home-no-load cases are included in the denominator. Twenty-four zero-load cases occurred for Raleigh on February 24 across both schedules and all policies; the remaining six were two each for Atlanta on March 3 and March 24 and Raleigh on March 10. A missing-board event and a feasible-load shortage are distinct conditions; no recorded no-board event does not prove that the archive completely represents available freight.

Table 2 gives per-policy means and medians of the individual terminal-start annualized projections. Every 6/1 row has n = 50 and every 12/2 row has n = 45. Values are rounded to the nearest dollar. These are descriptive in-sample results, with negative values retained.

Selector	6/1 mean	6/1 median	12/2 mean	12/2 median
Contribution per time: prefer two loads	-20,817	-17,993	7,441	9,909

Selector	6/1 mean	6/1 median	12/2 mean	12/2 median
Contribution per time: one load	-27,293	-27,841	1,780	2,054
Contribution per time: best one or two loads	-23,187	-16,556	6,617	7,606
Two loads: include return-home cost and time every time	236	1,239	6,230	6,185
Two loads: include return-home cost and time in final 4 days	-6,488	-1,341	13,630	22,194
Highest load revenue	-79,031	-74,552	-49,507	-48,692
Highest revenue per loaded mile	-30,416	-33,599	-4,424	-440
Shortest empty drive to pickup	-61,480	-59,851	-39,066	-34,391
Highest contribution on one load	-34,137	-39,129	-3,708	-3,935
Highest contribution per total load mile	-29,773	-29,489	-839	-789
Meet dispatch thresholds, then highest contribution	-41,595	-43,699	-17,937	-19,079
Contribution plus average destination value	-33,905	-39,129	-3,913	-3,903

The highest observed mean differs by schedule: always-on return for 6/1 and four-day return for 12/2. The gross-revenue and minimum-deadhead rules had substantially lower means in these data. This does not establish why: a selector changes the selected sequence, elapsed time, position, feasible future choices, and return movement simultaneously. It also does not mean that such rules cannot succeed under other market or cost conditions.

Matched differences help separate policy variation from the terminal-start mix. Table 3 reports policy contribution minus pair-preferred open contribution on the same terminal and start. The four-day 12/2 median is close to zero despite a positive mean, indicating that the average improvement is not representative of a uniformly large improvement in every scenario. These are paired descriptive differences, not confidence intervals or causal estimates for a new market.

Schedule and treatment	Matched cases	Wins / ties / losses	Mean difference per cycle	Median difference per cycle	Mean annualized difference
6/1: Two loads, return-home cost and time every time	50	37 / 2 / 11	403.76	396.01	21,053

Schedule and treatment	Matched cases	Wins / ties / losses	Mean difference per cycle	Median difference per cycle	Mean annualized difference
6/1: Two loads, return-home cost and time in final 4 days	50	24 / 25 / 1	274.80	0.00	14,329
12/2: Two loads, return-home cost and time every time	45	20 / 1 / 24	-46.46	-5.80	-1,211
12/2: Two loads, return-home cost and time in final 4 days	45	23 / 21 / 1	237.39	7.68	6,189

Wins, ties, and losses count terminal-start cases where the treatment's annualized result is respectively above, equal to, or below the open baseline. The four-day policy tied the baseline in 25 of 50 6/1 cases and 21 of 45 12/2 cases; these ties help explain its small median differences.

Figure 1 shows cumulative means of individual annualized results through each start week; Figure 2 shows the weekly median across the five terminals. The curves are descriptive summaries of overlapping historical scenarios, not a sequential deployment or a confidence band. No selector produced a positive weekly mean or median in every tested week. For 12/2, always-on return and four-day return each had a positive weekly mean in eight of nine weeks; four-day return also had positive weekly medians in eight of nine. For 6/1, each return-aware policy had positive weekly means in five of ten weeks. Figures use common dollar scales within each chart and include zero-load cases.

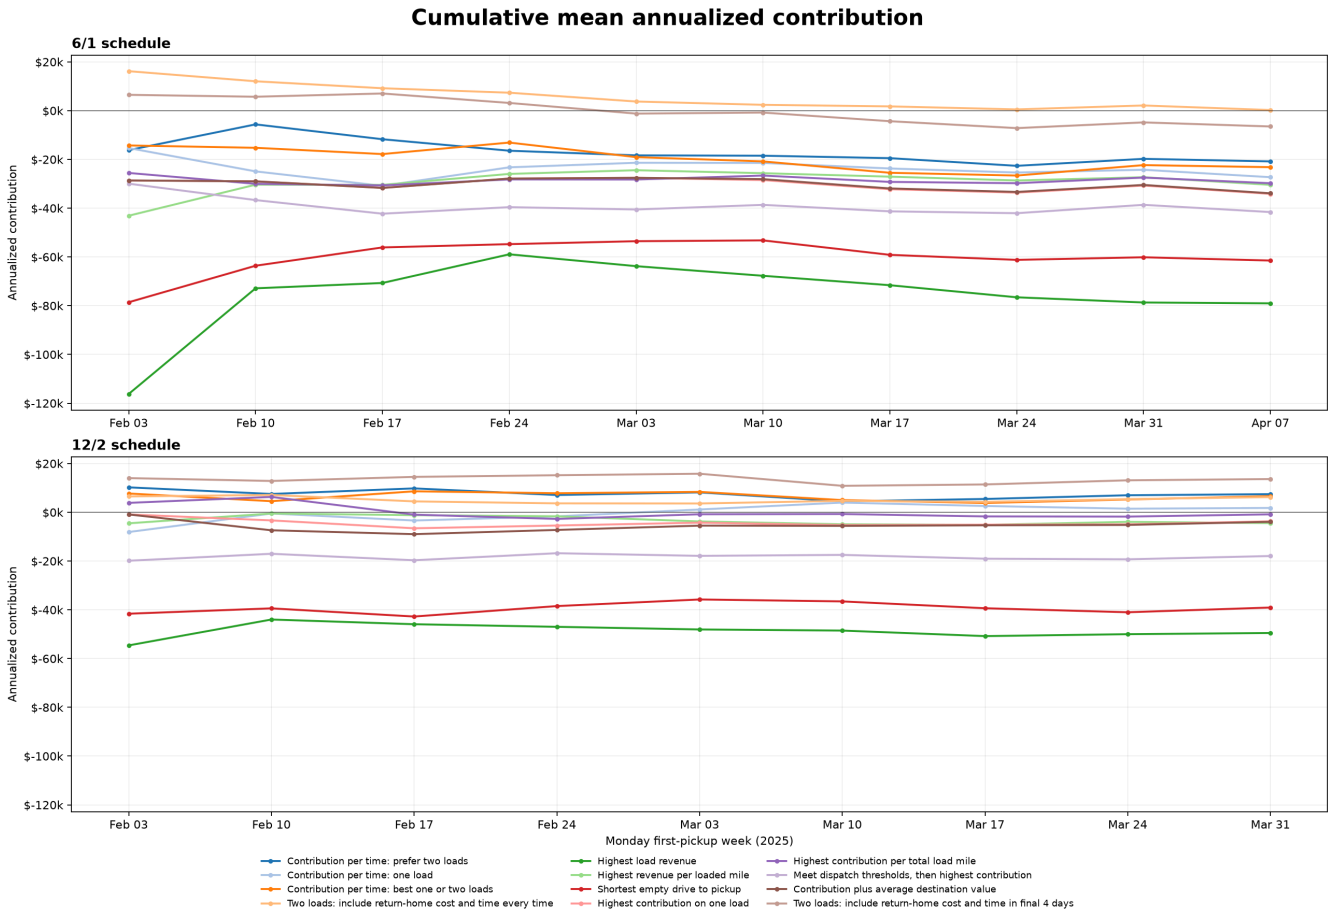


Figure 1. Cumulative mean annualized contribution for each schedule and selector. Each week contributes five terminal outcomes.

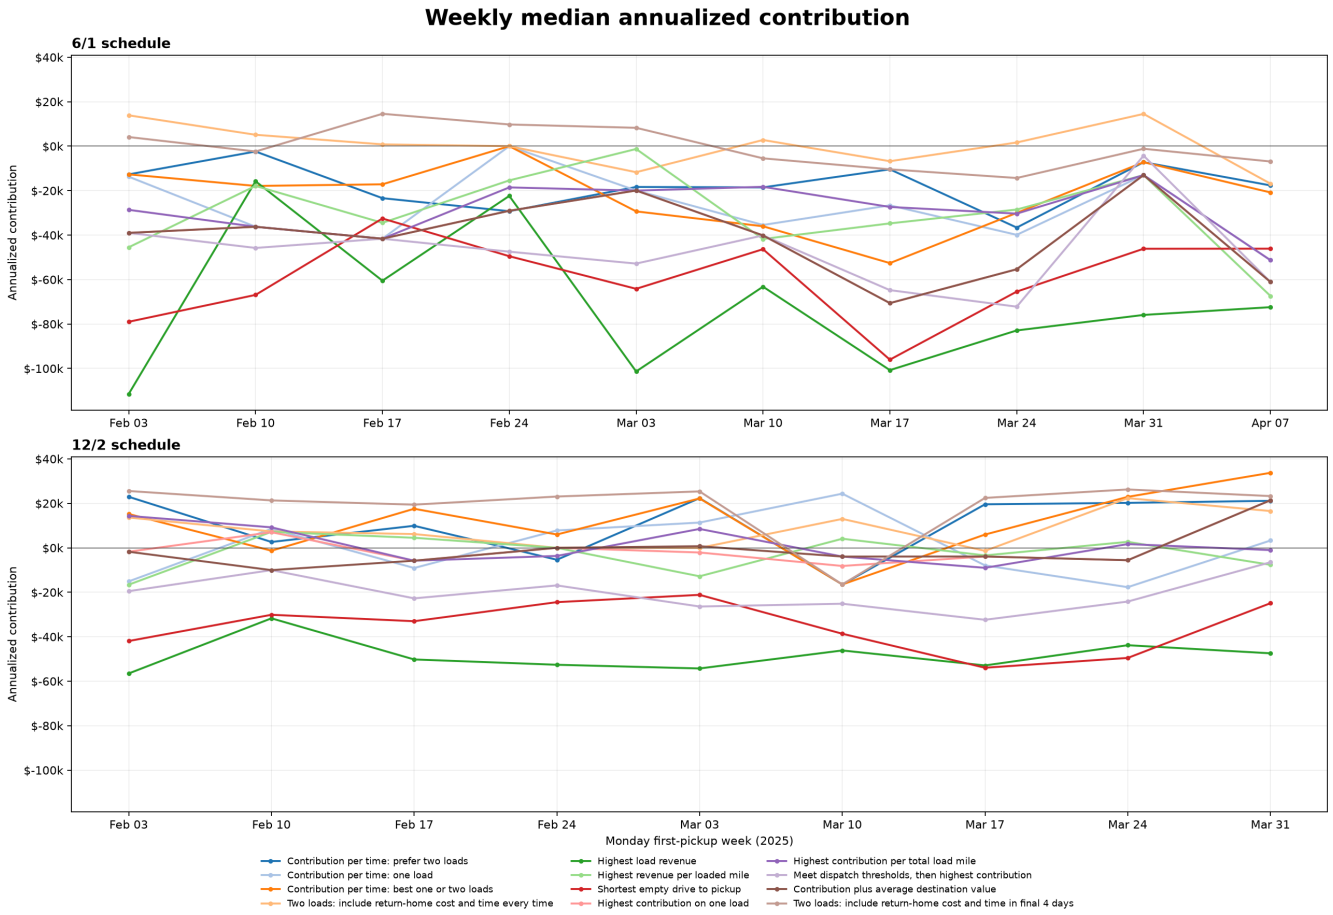


Figure 2. Median annualized contribution across five terminals for each selector and start week.

The terminal-level means across all 12 policies range from -$52,810 for Atlanta to -$8,115 for Columbus under 6/1, and from -$15,587 for Atlanta to $2,787 for Columbus under 12/2. These averages pool deliberately different policies, including weak baselines, and should not be interpreted as an estimate for a carrier operating a single chosen strategy. They nevertheless illustrate substantial geography-dependent variation within the sampled archive.

5.1. Gross revenue, activity, and public context

Gross revenue is the sum of posted prices of completed loads before modeled costs, commission, or an owner-operator split. The price field has not been reconciled to fuel surcharge and accessorial treatment; it should not be assumed to have the same billing boundary as a public carrier's reported revenue. Table 4 annualizes mean simulated cycle gross revenue, modeled costs, load counts, and miles using the same fixed calendar denominators as contribution. The

table includes valid zero-load cases. The 6/1 and 12/2 groups contain 50 and 45 cases per policy respectively; their date ranges differ by April 7. Annualized figures are repeated-cycle scenarios, not measured annual truck results.

Schedule and policy	Gross $/year	Modeled costs $/year	Contribution $/year	Loads/year	Loaded / pickup / home miles per year
6/1: prefer two loads, open	201,807	222,625	-20,817	181	77,505 / 10,573 / 29,266
6/1: two loads, always include home	158,101	157,865	236	157	69,116 / 8,619 / 5,476
6/1: two loads, home in final four days	198,501	204,989	-6,488	184	80,909 / 11,274 / 15,866
12/2: prefer two loads, open	177,373	169,932	7,441	162	66,919 / 9,584 / 13,067
12/2: two loads, always include home	129,832	123,602	6,230	126	54,819 / 7,199 / 3,132
12/2: two loads, home in final four days	177,548	163,918	13,630	160	68,952 / 9,567 / 7,881

Table 4. Annualized scenario means, rounded individually; mileage components are paid loaded, pickup deadhead, and recorded final homeward miles. Costs combine modeled completed-load and homeward costs. The accompanying aggregate comparison CSV contains all 12 selectors and separates load and homeward costs. Totals may differ by rounding.

The gross revenue and mileage figures explain why a higher revenue result need not yield higher modeled contribution. For example, the 6/1 open-pair policy produces more projected revenue than the always-home-valued policy but also incurs more modeled costs and homeward miles. These are whole-policy outcomes; the table alone cannot attribute their difference solely to the return leg. Across all 1,140 selected checkpoints, completed-load revenue matched the replay ledger and revenue less modeled load and home costs matched the reported contribution within $0.06 per cycle. This is an internal accounting check, not validation of the posted prices, reconstructed mileage, or assumptions about availability.

For scale, Knight-Swift reported $199,897 of 2025 truckload revenue per average tractor, excluding fuel surcharge and intersegment transactions, and 83,650 total miles per average tractor [8]. Werner reported $4,737 per tractor-week for one-way truckload service, net of fuel surcharge; multiplying by 52 gives a calculated $246,324 annual equivalent [7]. Their fleet results

cover actual full-year operations with service and equipment mixes unlike the simulated sample. Landstar reported $2,260 per BCO load excluding separately billed BCO fuel surcharges [11]. Adding the reported $228.9 million of separately billed BCO surcharges divided by 798,050 BCO loads gives approximately $2,547 per load before the BCO split [12]. By contrast, the 12/2 open-pair replay's ratio of aggregate completed-load revenue to completed loads is $1,093. The figures are displayed as context, not a performance gap: price and fuel treatment, length of haul, geography, equipment, contract mix, and service definitions are not aligned.

Credible external comparison requires tracing the archive's price and surcharge fields; matching loaded, pickup, return, and other empty-mile definitions; aligning driver compensation, contractor split, overhead, and other costs; and comparing common service types and market periods. In particular, Landstar's revenue before the BCO split, a BCO's receipts after the split, and this paper's modeled carrier gross are different economic entities' measures. Applying an owner-operator split after the fact also does not show how a selector would perform if re-optimized under the owner-operator's economics.

6. Schedule interpretation and dispatcher implications

On the 45 common terminal-start cases, pair-preferred open averaged -$19,764 annualized under 6/1 and $7,441 under 12/2, a paired average difference of $27,205 in favor of 12/2; its paired median difference was $25,264. These are separate declared-cycle projections. The analysis does not simulate two state-connected 6/1 cycles against one 12/2 cycle over the same 14 days, and cannot isolate the effect of avoiding an intermediate return or establish the value of the extra time away from home. Home arrival, opportunities, route choices, and score trajectories differ across arms. The April 7 6/1 observations do not enter this matched schedule comparison.

For a dispatcher, the findings support presenting several feasible choices with their modeled tradeoffs rather than relying on a single rate or distance measure. A practical display could show incremental contribution, time until availability, pickup deadhead, projected home cost and feasibility, the board capture time, and a next-best alternative. The home-valued policy may prevent an attractive outbound choice from creating an expensive return; it may also reject a useful load too early. Four-day and always-on rules provide development evidence for that tradeoff, not a universal start-home threshold. A destination opportunity adjustment and a two-load bundle solve different problems: the latter scores a specific candidate sequence, whereas a destination prior estimates value after the final delivery. The legacy prior here is not evidence for a top-5% destination model.

7. Limitations and future validation

The archive is a partial collection of board snapshots, not a complete record of market supply, accepted transactions, rates actually paid, or live load availability. Reuse of an earlier board and retrospective knowledge of missing capture days can materially affect routes. The posted-price, routing, service-time, cost, and HOS assumptions require operational calibration; a feasible simulated home arrival is not a verified real-world execution. Candidate construction differs between graph and simpler selectors, and the legacy backhaul ranking has not been fully independently reconstructed. Modeling return-home feasibility does not model the driver's full lived schedule or preferences.

All policies, terminal locations, and neighboring start dates share archived freight; 12/2 cycles overlap in calendar time. The 1,140 policy outcomes are not independent replications. Policy development and selection used the same period, and there is no held-out or multiseason validation. Means, medians, and positive-week counts are descriptive; no statistical significance or national ranking is asserted. A positive projected contribution also does not establish an adequate return after omitted overhead, owner labor, financing, or other business-specific costs.

Next, test the selected rules without retuning on later and geographically broader data with verified capture timing, and compare against shadow-mode dispatcher decisions and confirmed offers. Record booking attempts, prices, rejection or unavailability, actual appointments, driver and route outcomes, and complete costs. Compare connected two-cycle 6/1 trajectories with one 12/2 trajectory over a common 14-day horizon without inventing a reset, and quantify the intermediate home leg and driver time separately. Further sensitivity tests should vary lead time for home valuation, candidate availability, board-gap treatment, and cost profiles, and should separately evaluate owner-operator economics. Report matched scenario distributions, failures, computational latency, and uncertainty using a design that accounts for shared weeks and geography. Disaggregated loaded, pickup-deadhead and homeward miles, wait time, load counts, and representative matched route traces are needed to investigate the mechanisms behind these outcomes; they have not been independently synthesized in this paper.

8. Conclusion

In this historical development replay, incorporating homeward value was associated with higher average contribution than the open pair policy for 6/1 and, with a four-day trigger, for 12/2. The corresponding 12/2 matched median gain was small, and no selector had positive weekly means

or medians in every week. The study shows why full-cycle contribution, home-time constraints, and adverse scenarios matter when comparing dispatch rules. Its partial board coverage, in-sample design, and modeled operations limit the findings to hypotheses for prospective testing.

Reproducibility and provenance

Working paper, revised 10 November 2026. The combined development report records each selected outcome, attempts, source identity checks, and matched comparisons. The archive SHA-256 is 03d2ec61c6b740df73d5ce01742897d8a6af482023fe7c6c609fc1cbede70f22; the shared input identity is 8ca62c323c91eb4a3f047c86c288e020fc23482a9529470190badbe60884e6d8. Two reviewed replay source identities differ by study-date allowlist changes. The local research summary is `freight-all-mondays-results-20261110.md` in the recommendation-engine benchmarks at commit 2a896630; the plan completion record is commit f3249617. Detailed artifacts remain in the local shared workspace and are not a public dataset; the material available here is insufficient for independent external replication without access to that archive and its manifests. Eight focused coordinator and reporting tests passed. Independent confirmation of candidate ranking for every legacy alternative and access to a public reproducibility package remain future work.

References

[1] Miller, J., Nie, Y., and Liu, X. (2020). Hyperpath Truck Routing in an Online Freight Exchange Platform. Transportation Science 54(6), 1676–1696. https://doi.org/10.1287/trsc.2020.0989.

[2] Dong, Y., Bai, Y., Olinick, E., and Yu, A. J. (2022). The Backhaul Profit Maximization Problem: Optimization Models and Solution Procedures. INFORMS Journal on Optimization 4(4), 373–402. https://doi.org/10.1287/ijoo.2022.0071.

[3] Dumouchelle, J., Frejinger, E., and Lodi, A. (2023). Reinforcement Learning for Freight Booking Control Problems. arXiv preprint, version 3. https://arxiv.org/abs/2102.00092.

[4] Chandrasekaran, A. (2026). Latency-Aware Bid Acceptance under Operational Feasibility: A Public Benchmark with Hindsight Ceilings. arXiv preprint. https://arxiv.org/abs/2607.07343.

[5] Ojha, R., Chen, W., Zhang, H., Khir, R., Erera, A., and Van Hentenryck, P. (2024). Optimization-based Learning for Dynamic Load Planning in Trucking Service Networks. arXiv preprint, version 2. https://arxiv.org/abs/2307.04050.

[6] American Trucking Associations. Economics and Industry Data, 2024 trucking revenue and June 2025 registered-carrier fleet-size summary. https://www.trucking.org/economics-and-industry-data. Accessed 10 November 2026.

[7] Werner Enterprises. Form 10-K for 2025, Truckload Transportation Services segment. https://www.sec.gov/Archives/edgar/data/793074/000079307426000071/wern-20251231.htm.

[8] Knight-Swift Transportation Holdings. Form 10-K for 2025, truckload segment. https://www.sec.gov/Archives/edgar/data/1492691/000149269126000016/knx-20251231.htm.

[9] Marten Transport. Form 10-K for 2025, truckload and brokerage segments. https://www.sec.gov/Archives/edgar/data/799167/000143774926005971/mrtn20251231_10k.htm.

[10] ATBS. How Did Owner-Operators Perform in 2025? https://www.atbs.com/post/how-did-owner-operators-perform. Accessed 10 November 2026.

[11] Landstar System. Form 10-K for fiscal 2025, BCO independent-contractor operating statistics and non-GAAP variable contribution reconciliation. https://www.sec.gov/Archives/edgar/data/853816/000119312526064756/d66072d10k.htm.

[12] Landstar System. Fourth-quarter and year-end 2025 earnings release, BCO fuel-surcharge disclosure. https://investor.landstar.com/static-files/affadf70-892d-446d-86ea-41fbe00cfba5.

Glossary

Freight and dispatch terms

- Backhaul: A load or trip taken after an outbound delivery, often in the direction of home or another desired market. A backhaul need not return the truck all the way home.
- Booking: Committing to a load. In this study, booking is simulated from a captured load board; it does not confirm that a broker accepted the booking.
- Broker: An intermediary that arranges transportation between a shipper and a carrier.

- Carrier: The trucking business responsible for moving freight. This study uses a modeled carrier profile with employee driver pay and benefits.
- Deadhead: Empty travel. Pickup deadhead is the drive from the truck's current position to a load's pickup; homeward deadhead is an empty trip toward home.
- Dispatcher: The person or system that evaluates available freight and helps assign loads and plan a truck's movements.
- Hours of service (HOS): Rules governing a commercial driver's driving, on-duty, and rest time. The study uses a simulated HOS model, not a legal compliance determination.
- Load board: A marketplace or feed of freight offers. A board snapshot or capture is a record of offers observed at a particular time, not a complete record of available or booked freight.
- Loaded miles: Miles traveled while hauling a load, as represented by its reported loaded mileage in this study.
- Over the road (OTR): Truck operations involving trips away from a home terminal. The study uses a 250-loaded-mile minimum as an OTR proxy, not a universal definition.
- Posted price: The amount shown for a load offer. The replay treats this as modeled revenue without validating negotiation, acceptance, or actual payment.
- Revenue per mile (RPM): Revenue divided by a stated mileage measure. Loaded RPM uses loaded miles; all-mile RPM in this study uses pickup deadhead plus loaded miles.
- Terminal: The truck's designated home location for a replay, including the location used to assess its return-home deadline.
- Truckload: Freight moved as a load assigned to a truck or trailer. This paper studies modeled spot-market truckload choices, not coordinated assignment of an entire fleet.

Study and measurement terms

- 6/1 and 12/2: Declared cycles of six days out and one day home, or twelve days out and two days home. The study projects each simulated cycle separately; it does not connect two 6/1 cycles into a matched 14-day trajectory.
- Annualized contribution: A fixed-cycle projection calculated as cycle contribution × 365/7 for 6/1 or × 365/14 for 12/2, including when the driver returns early. It is not observed annual earnings.
- Contribution: Modeled load or cycle revenue less the costs included in the study's economics profile. It is not necessarily accounting profit after all business expenses.
- Contribution per time: The local ranking measure that relates a candidate load or bundle's modeled contribution to its elapsed committed time. It differs from the fixed-cycle annualized result reported for a completed replay.

- Cycle: One simulated declared work-and-home period. Cycle contribution includes modeled load and return-home costs under the study's accounting rules.
- Destination value or prior: An estimate of future opportunity associated with the place a load delivers. The legacy selector uses an average-based prior; it is not the separately considered top-5% destination measure.
- Eligibility screen: A rule that removes candidates before selection. The study requires a positive posted price, at least 250 reported loaded miles, and a mileage-consistency check, in addition to implemented equipment, timing, and feasibility rules.
- Home deadline: The simulated time by which the truck must return to its designated terminal. A modeled on-time return is not proof of real-world execution.
- Home-valued selection: Ranking candidates using an estimated immediate return-home cost and time after the candidate loads. The resulting estimate is part of the selection score, not a guarantee that the driver will take that exact return at that moment.
- Matched comparison: A comparison between policies using the same terminal, start date, and schedule. A win, tie, or loss indicates whether the treatment's outcome was above, equal to, or below the baseline for that case.
- Missing-dispatch fallback: The replay's use of an earlier available board when the expected dispatch-day capture is missing. It relies on retrospective knowledge of capture gaps and does not establish that an earlier offer remained bookable.
- No-load outcome: A replay that returns or remains home without completing a load. Valid on-time no-load cases are retained as zero contribution in the study's averages.
- Open selection: A candidate-ranking objective that does not add an immediate home trip to the candidate's score, while still requiring a feasible home return. The completed cycle's accounting can still include homeward cost.
- Replay: A simulation that makes decisions using recorded historical board observations and modeled truck state, timing, costs, and feasibility. It is not a record of actual dispatch decisions or realized earnings.
- Selector or policy: A rule or algorithm for choosing among eligible candidate loads or bundles. Selector scores guide individual decisions; cycle contribution evaluates the resulting simulated route.
- Two-load bundle: A candidate sequence of up to two loads considered together from one revealed board. Pair-preferred policies choose a feasible pair when available and otherwise fall back to one load; the replay does not prove both offers were simultaneously bookable.

- Weekly mean and median: For a selector and Monday start, the mean is the average of the five terminal outcomes, and the median is the middle of those five outcomes. A positive-week count includes only values strictly above zero.

Selector names

- Contribution per time: prefer two loads: Prefers a feasible two-load bundle and ranks it by modeled contribution relative to elapsed time, without adding an immediate return-home trip to the score; falls back to one load.
- Contribution per time: one load: Ranks individual loads by modeled contribution relative to elapsed time, without adding an immediate return-home trip to the score.
- Contribution per time: best one or two loads: Compares eligible one-load and two-load candidates using modeled contribution relative to elapsed time, without adding an immediate return-home trip to the score.
- Two loads: include return-home cost and time every time: Prefers a feasible pair and uses an estimated immediate return-home cost and time in the ranking objective; falls back to one load.
- Two loads: include return-home cost and time in final 4 days: Prefers a feasible pair and switches to home-valued ranking when projected availability is within four elapsed days of the deadline; the initial booking uses the open objective.
- Highest load revenue: Selects for the highest posted revenue on an eligible load, without dividing by miles or subtracting modeled cost in the score.
- Highest revenue per loaded mile: Selects for the highest posted revenue divided by loaded miles, excluding pickup deadhead from the denominator.
- Shortest empty drive to pickup: Selects for the least pickup deadhead, regardless of the load's price in the ranking score.
- Highest contribution on one load: Selects the eligible load with the highest modeled incremental contribution in dollars, without dividing by its elapsed time.
- Highest contribution per total load mile: Ranks by posted price divided by pickup deadhead plus loaded miles. Under the study's constant per-mile cost and zero commission, this gives the same ranking as modeled contribution per those miles; later repositioning and return-home miles are excluded from the score.
- Meet dispatch thresholds, then highest contribution: Rejects loads below $1.50 revenue per loaded mile or below zero modeled incremental contribution, or above 150 pickup-deadhead miles or 36 hours of pickup wait; then favors the highest incremental contribution among the remaining loads.

- Contribution plus average destination value: Adds a legacy average-based destination opportunity prior to modeled incremental contribution. Its candidate ranking has not been fully independently reconstructed.